

Modeling and Control of Aggregated Heterogeneous Thermostatically Controlled Loads for Ancillary Services

Stephan Koch
ETH Zurich
Power Systems Laboratory
koch@eeh.ee.ethz.ch

Johanna L. Mathieu
University of California, Berkeley
Dept. of Mechanical Engineering
jmathieu@berkeley.edu

Duncan S. Callaway
University of California, Berkeley
Energy and Resources Group
dcal@berkeley.edu

Abstract - This paper presents a novel modeling and control approach for the aggregation of large numbers of heterogeneous thermostatically controlled loads, such as refrigerators, electric water heaters, and air conditioners, and their usage for Demand Response. Unlike traditional Demand Response methods that act on time scales of hours, this approach is able to provide short-term (e.g., second-to-second) ancillary services, such as balancing and frequency control. A statistical modeling approach based on Markov Chains is used to describe the evolution of probability mass in a temperature state space. The Markov state transition matrix is identified using state information from the population of thermostatically controlled loads. A predictive controller is used to control the aggregate population of loads such that it tracks a signal. A simulation example shows the applicability of the approach to realistic systems, and includes a comparison of control performance depending on available state information and controller parameterization.

Keywords - Demand Response, Thermostatically Controlled Loads, Markov Chains, Predictive Control, Aggregation Strategies, Ancillary Services.

Notation

Symbol	Unit	Meaning
a	[-]	TCL parameter
C	[kWh/°C]	Thermal capacitance
h	[hr]	Time step
i	[-]	TCL index
m	[-]	ON/OFF switching variable
N	[-]	Number
R	[°C/kW]	Thermal resistance
P	[W]	Power
w	[°C]	Noise process
\mathbf{u}_{bin}	[-]	Control input vector
\mathbf{v}	[W]	Measurement noise vector
\mathbf{x}	[-]	State bin vector
δ	[°C]	Temperature dead-band width
θ	[°C]	Temperature

Subscript	Meaning	Subscript	Meaning
a	ambient	rate	TCL rating
bin	state bins	set	setpoint
g	gain	total	whole population
p	power	track	tracking
pred	prediction steps		

1 Introduction

THE growing capacity of renewable electricity generators with variable and uncertain production patterns has resulted in an acute need for increased system flexibility [1]. At the same time, the expansion of grid cyberinfrastructure (i.e. the “smart grid”) is creating a significant opportunity to manage the output of large aggregations of electric loads with a level of precision comparable to (or beyond) supply-side capacity, without significant reductions in the end-use performance of the loads. A large fraction of demand on the grid lies in thermo-

statically controlled loads (TCLs) whose local controllers maintain temperatures within a dead-band using hysteretic (ON/OFF) control. Because these loads have naturally occurring “slack” [2] with respect to local temperature and time of operation, they represent a promising end-use category to engage in power system services.

Several recent papers have explored the potential to engage TCLs in power system services with detailed end-use load models. For example, [3] develops a state queueing model of TCLs for the purpose of understanding the potential response of TCLs to external inputs such as a time-varying electricity price. [4] examines a market-based scheme where loads switch early from ON to OFF or vice versa based on willingness to accept payment to do so, and [5] develops a system-level modeling scheme for incorporating TCLs into managing wind forecast error. In [6], aggregate TCL power consumption is manipulated by thermostat setpoint control, and a statistical mechanical model is solved to gain insight into the system dynamics. [7] extends this work to a discrete “state bin” modeling framework similar to [3], and the authors develop a bilinear controller to manipulate aggregate power consumption, again by temperature setpoint control.

The objective of this paper is to build on these previous studies to develop a general load modeling and control framework that is flexible with respect to parameter and state estimation and control design. At its core, the approach is to develop a linear time-invariant (LTI) representation of a population of TCLs. The modeling strategy is similar to [3, 7] in that we describe local device states in discrete temperature-related state bins and allow aggregated probability mass to move through these bins. In a manner similar to [7], we are interested in understanding whether system identification and observer design methods can be brought to bear on the aggregated TCL system.

The work we present in this paper is novel in several domains. First, the aggregated state space modeling framework we present explicitly models TCL heterogeneity. Second, we show that the aggregated model parameters can be determined either from parameters of the TCLs being controlled, or by observation of the temperature dynamics of some or all of the loads in the population. Third, we formulate the control problem in terms of a Model Predictive Control (MPC) framework that is highly flexible with respect to managing control objectives and constraints. Finally, anticipating that it will be difficult to measure in real time the state of every TCL being controlled, we explicitly examine the influence that incomplete TCL state information has on controller performance.

The paper proceeds as follows. In Section 2, we describe our modeling approach and derive an LTI system that captures the temperature dynamics of the TCL population. In Section 3, we present different options for information exchange between the central controller and the TCL population. Section 4 presents the MPC approach, which is tested in a case study in Section 5. In Section 6, we present ideas for further research and conclude.

2 Modeling Approach

2.1 Individual TCL Model

The individual TCL model is the same as that used in [6]. The parameter, a , which governs the thermal characteristics of each TCL, i , is defined as:

$$a_i = e^{-h/C_i R_i} \quad , \quad (1)$$

where C_i and R_i are the thermal capacitance and resistance of TCL i , and h is the time step.

The difference equation describing the temperature dynamics of TCL i is:

$$\theta_{i,t+1} = a_i \theta_{i,t} + (1 - a_i)(\theta_{a,i} - m_{i,t} \theta_{g,i}) + \omega_{i,t} \quad , \quad (2)$$

where θ is the internal temperature of the TCL, θ_a is the ambient temperature, m is a dimensionless discrete variable equal to 1 when the TCL is ON and 0 when it is OFF, and ω is a noise process. The ON temperature gain is:

$$\theta_{g,i} = \begin{cases} R_i P_{\text{rate},i} & \text{for cooling devices} \\ -R_i P_{\text{rate},i} & \text{for heating devices} \end{cases} \quad , \quad (3)$$

where $P_{\text{rate},i}$ is the TCL's rated power. $\theta_{g,i}$ is positive for cooling TCLs and negative for heating TCLs.

2.2 TCL Population Model

To simulate the behavior of a population of TCLs, we could aggregate thousands of single TCL models using (2); however, this would be computationally intensive and the aggregate system would not be in a form amenable to many control techniques. Instead, in this paper we will work with a discrete LTI system in state space form:

$$\mathbf{x}_{k+1} = \mathbf{A}\mathbf{x}_k + \mathbf{B}\mathbf{u}_k \quad (4)$$

$$\mathbf{y}_k = \mathbf{C}\mathbf{x}_k \quad , \quad (5)$$

which allows us to use a wide range of system analysis tools and advanced controls techniques.

Assume all TCLs in a population have the same temperature setpoint, θ_{set} , and temperature dead-band width, δ , (or normalize diverse dead-bands). Divide the dead-band into $N_{\text{bin}}/2$ temperature intervals. A TCL in a certain temperature interval can be either ON or OFF. Divide each temperature interval into two state bins, one for TCLs that are ON and one for TCLs that are OFF. This results in N_{bin} state bins. The state vector \mathbf{x} contains the number of TCLs in each state bin, or, if normalized by the total number of TCLs, the fraction of TCLs in each state bin, which is equivalent to probability mass in the infinite system limit. In the remainder of the paper, we will refer to

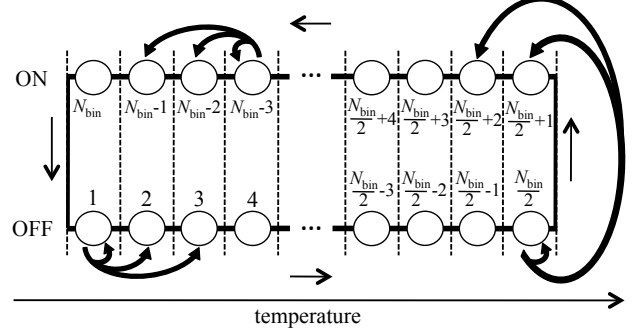


Figure 1: State bin transition model.

\mathbf{x} as a vector of probability mass. The \mathbf{A} -matrix can be thought of as a (transposed) Markov transition matrix describing the probability of TCLs moving from one state bin to the next. Figure 1 shows how the state bins map to the temperature dead-band.

We next present an analytical derivation of the \mathbf{A} -matrix, which we compare to an identified \mathbf{A} -matrix. We will discuss the structure of the \mathbf{B} and \mathbf{C} matrices and \mathbf{u} and \mathbf{y} vectors in subsequent sections.

2.2.1 Analytical Derivation of the \mathbf{A} -matrix

In this section, we will analytically derive the elements of the \mathbf{A} -matrix. The purpose of this exercise is to demonstrate that the modeling framework is rooted in basic probabilistic logic. In subsequent sections, we will estimate the \mathbf{A} -matrix with data generated by simulating (2).

We assume that all TCLs have the same resistance ($R_i = R$) and rated power ($P_{\text{rate},i} = P_{\text{rate}}$). This implies that all TCLs have the same temperature gain ($\theta_{g,i} = \theta_g$). We also assume that all TCLs experience the same ambient temperature ($\theta_{a,i} = \theta_a$), which is constant over time. Lastly, we ignore the noise process ω . Therefore, (2) becomes:

$$\theta_{i,t+1} = a_i \theta_{i,t} + (1 - a_i)(\theta_a - m_{i,t} \theta_g) \quad . \quad (6)$$

Take a TCL going from θ_{start} to θ_{end} in one time step:

$$\theta_{\text{end}} = a_i \theta_{\text{start}} + (1 - a_i)(\theta_a - m_{i,t} \theta_g) \quad . \quad (7)$$

Now consider a group of TCLs that are either all ON or all OFF ($m_{i,t} = m_t$). The probability of TCLs going from θ_{start} to θ_{end} is:

$$P(\theta_{\text{end}} | \theta_{\text{start}}) = P(a_i) \quad , \quad (8)$$

where we have assumed that $P(a_i)$ is independent of temperature. This probability can be computed by solving for a_i :

$$a_i = 1 - \frac{\theta_{\text{end}} - \theta_{\text{start}}}{\theta_a - \theta_{\text{start}} - m_t \theta_g} = \frac{\theta_a - \theta_{\text{end}} - m_t \theta_g}{\theta_a - \theta_{\text{start}} - m_t \theta_g} \quad . \quad (9)$$

Consider the same group of TCLs. The probability of TCLs going from θ_{start} to some θ_{end} , which is within a state bin ($\theta_n < \theta_{\text{end}} < \theta_{n+1}$), is:

$$\begin{aligned} P(\theta_n < \theta_{\text{end}} < \theta_{n+1} | \theta_{\text{start}}) & \quad (10) \\ & = P(a_1 < a < a_2) = \int_{a_1}^{a_2} p(a) da, \end{aligned}$$

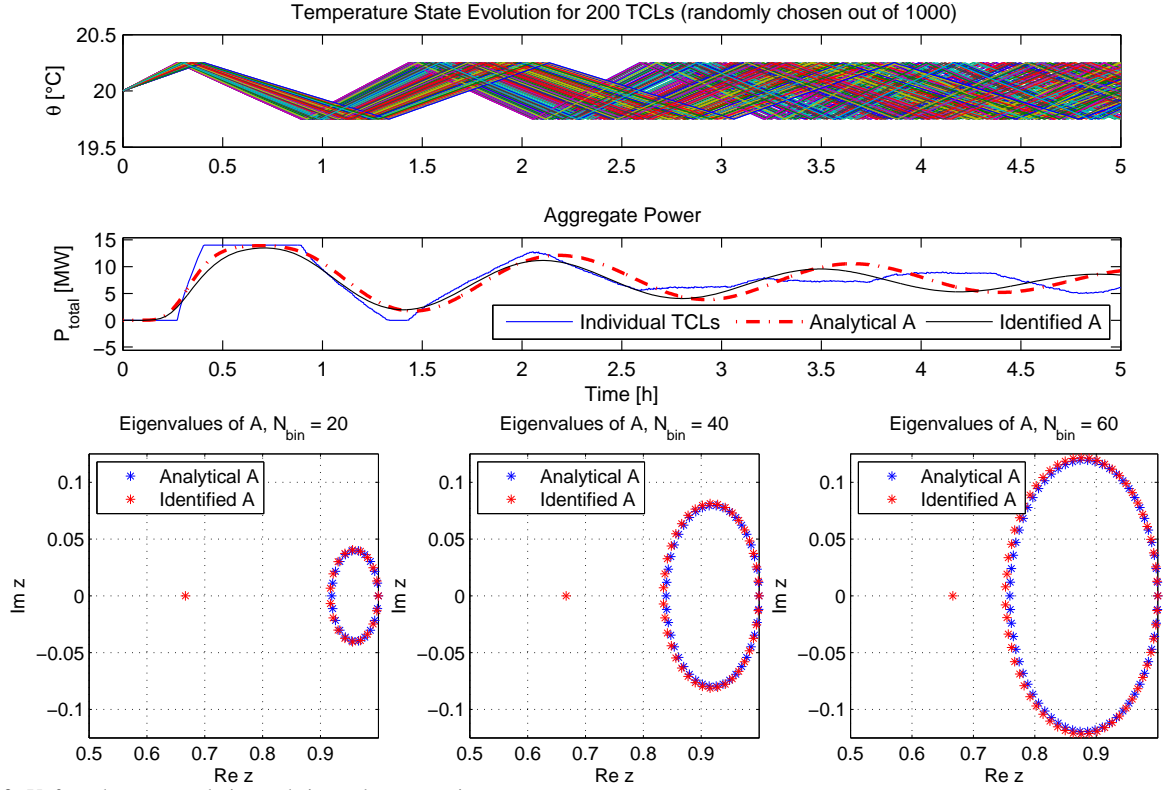


Figure 2: Unforced system evolution and eigenvalue comparison.

where

$$a_1 = \frac{\theta_a - \theta_1 - m_t \theta_g}{\theta_a - \theta_{\text{start}} - m_t \theta_g} \quad (11)$$

$$a_2 = \frac{\theta_a - \theta_2 - m_t \theta_g}{\theta_a - \theta_{\text{start}} - m_t \theta_g} \quad (12)$$

Here, $\theta_1 = \theta_n$ and $\theta_2 = \theta_{n+1}$ if the TCL is traversing from low temperatures to high temperatures (cooling/ON or heating/OFF), or $\theta_1 = \theta_{n+1}$ and $\theta_2 = \theta_n$ if the TCL is traversing from high temperatures to low temperatures (cooling/OFF or heating/ON).

Again, consider the same group of TCLs. Assume all state bins ($\theta_m < \theta < \theta_{m+1}$) contain TCLs that are uniformly distributed over the temperature range of the bin AND the parameter distribution $p(a_i)$ is independent of temperature. Then the probability of TCLs going from $\theta_m < \theta_{\text{start}} < \theta_{m+1}$ to $\theta_n < \theta_{\text{end}} < \theta_{n+1}$ is:

$$\begin{aligned} & P(\theta_n < \theta_{\text{end}} < \theta_{n+1} | \theta_m < \theta_{\text{start}} < \theta_{m+1}) \\ &= \int_{\theta_m}^{\theta_{m+1}} \int_{a_1}^{a_2} p(a) da d\theta_{\text{start}}, \quad (13) \end{aligned}$$

with a_1 and a_2 defined as in (11) and (12). We evaluate (13) for each combination of starting and ending bins to derive the \mathbf{A} -matrix.

2.2.2 Comparison of Analytically-derived \mathbf{A} -matrix to Identified Matrix

We compare the analytically-derived \mathbf{A} -matrix to an identified matrix obtained through a Markov chain construction process. The application of this technique implies that state information from the TCL population (or a

Parameter	Meaning	Value
θ_{set}	temperature setpoint	20°C
δ	temperature dead-band width	0.5°C
θ_a	ambient temperature	32°C
R	thermal resistance	2 °C/kW
C	thermal capacitance	8-12 kWh/°C
P_{rate}	power	14,000 W
h	time step	10 s

Table 1: Simulation and TCL parameters, adapted from [6].

subset of the population) is available for system identification purposes.

For this comparison, we consider the case where the TCL parameter, a , is uniformly distributed between a_{min} and a_{max} :

$$p(a) = \begin{cases} \frac{1}{a_{\text{max}} - a_{\text{min}}} & \text{if } a_{\text{min}} \leq a \leq a_{\text{max}} \\ 0 & \text{otherwise} \end{cases} \quad (14)$$

The simulation and TCL parameters used in our analysis are listed in Table 1, and a_{min} and a_{max} are computed from these parameters.

Analytically-derived matrix: For each combination of starting and ending bins, we evaluate (13) numerically to generate the analytically-derived \mathbf{A} -matrix. ON/OFF switching is dealt with through the addition of extra state bins, which are outside of the dead-band. In each time step, probability mass is moved from starting bins to ending bins, without regard to the dead-band. Before the start of the next time step, all probability mass that ends up in state bins outside of the dead-band is switched (ON to OFF, or vice versa).

Markov-identified matrix: We simulate 1000 individual TCLs and use the MATLAB function `hmmestimate` to identify the Markov Transition Matrix. We could also

construct the matrix by counting the number of TCL transitions from each combination of starting and ending bins over many time steps. Note that the Markov Transition Matrix has to be transposed to be equivalent to the \mathbf{A} -matrix from (4).

Model comparison: The dynamic behaviors of the analytically-derived and Markov-identified models were compared to that of a simulated population of 1000 individual TCLs. The top two plots in Figure 2 show the evolution of the unforced systems ($N_{\text{bin}} = 40$ for the analytical and identified \mathbf{A} matrices). The bottom three plots compare the eigenvalues of the analytically-derived and Markov-identified \mathbf{A} -matrices, given different numbers of state bins, here simulated with a time step of 10 s.

For the time-domain simulations, all TCLs are started in a single state bin, which leads to an oscillatory decay in aggregate power. The models produce similar results for the first 2-3 oscillations. After that, the disagreement between the three models increases. This implies that control actions that require long forecasts may not perform well, but that control actions with reasonably short forecast requirements (e.g. less than one hour) may not suffer from plant-model mismatch problems. If control actions are applied continuously to the system (“tight” control), this mismatch problem will likely be negligible.

The oscillatory decay is related to both TCL parameter heterogeneity and the number of state bins used in the state bin transition model. In the individual TCL model, homogeneous populations exhibit undamped oscillations, and these oscillations decay with increasing parameter heterogeneity as TCLs spread out over the temperature state space. In the bin transition model, there is a negative correlation between the number of bins and the damping of the oscillation: with increasing amount of bins, the oscillation becomes increasingly persistent. This finding is consistent with [7], where a large number of bins (200) was chosen in order to yield a non-decaying oscillation corresponding to a homogeneous system. Interestingly, the bin transition model shows a non-decaying behavior when large numbers of bins are used, even if the model has been constructed from a heterogeneous parameter set. Thus, for heterogeneous populations, the number of bins can be seen as a tuning variable which can be adapted to achieve a good matching of the decay behavior. However, as will be seen in Section 5, the number of bins does not have a large influence on a tightly controlled system.

The mismatch between the analytically-derived model and the simulated population of TCLs is due, at least in part, to the assumption made in Section 2.2.1 that the parameter distribution of TCLs is the same in each state bin. For the simulated population of TCLs, the parameter distribution of TCLs in each state is different and depends on initial condition and simulation time.

The mismatch between the analytically-derived and the Markov-identified \mathbf{A} -matrices can likely be attributed to the way ON/OFF transitions are handled. As described previously, in the analytical system, we have added extra bins outside of the temperature dead-band. Alternatively,

the identified matrix has only one bin at each corner of the temperature dead-band, which takes all probability mass that crosses the dead-band during each time step.

We conclude that the analytically-derived and Markov-identified models are relatively consistent with the simulated population of TCLs, though in both cases some error results from the matrix construction methodology. We use the (transposed) Markov matrix in our subsequent simulations.

3 Information Transfer

3.1 General Considerations

In this paper, we will not identify specific communications platforms or protocols. Instead, we will speak more abstractly about the types of information that might be exchanged. This work is based on a “direct load control” philosophy, meaning that a central controller collects information from the TCL population and transmits a control signal to which the TCLs react. Given the state bin transition model, there are many degrees of freedom for control and information transfer between the TCL population and the central controller including:

- A) the nature of the control signal transmitted from the central controller to the TCLs,
- B) the methodology for addressing the entire population or only a certain subset, and
- C) the information that the TCLs transmit back to the controller.

In Table 2, we present a number of options for these three degrees of freedom. Clearly, our choices of options must correspond to capabilities of the actual system. In the following subsections, we explain the subset of options we have implemented. Future work will explore the other options.

3.2 Implemented Approach

3.2.1 Controller \rightarrow TCL Population

We have implemented A3/B3 according to Table 2. Specifically, we toggle ON/OFF TCLs in certain bins, which corresponds to moving TCLs from ON bins to OFF bins of equal temperature, or vice versa. The control input vector \mathbf{u}_{bin} only has half the number of independent elements as the bin state vector since for every bin that loses probability mass there will be another equal-temperature state bin that gains. Considering the sequence of bins in Figure 1, the control input \mathbf{u}_{bin} enters the system via the matrix \mathbf{B} :

$$\mathbf{B} = \begin{bmatrix} -1 & & 0 \\ & \ddots & \\ 0 & & -1 \\ 0 & & 1 \\ & \ddots & \\ 1 & & 0 \end{bmatrix}. \quad (15)$$

Options	Feasible with State Bin Transition Model?	Implemented in this paper?
A. Control signal from central controller to TCLs		
A1. Shifting, expanding, or contracting the temperature dead-band, as in [6, 7]	yes, through definition of extra bins	no
A2. Blocking (switching OFF) TCLs at certain times	yes, through definition of extra bins	no
A3. Toggling ON/OFF certain TCLs, as in [4]	yes	yes
B. Targeted TCLs		
B1. Selective addressing of individual TCLs	no, unless the current bin of each TCL is known	no
B2. Broadcast same control to all TCLs with uniform response	yes	no
B3. Broadcast same control to all TCLs with differing responses	yes	yes
C. Information from TCLs to central controller		
C1. Full state information of the entire population	yes, through system ID	yes
C2. Full state information of a subset of the population	yes, through system ID and could include an observer	yes
C3. No state information, only aggregate power measured	yes, but may be difficult	no
C4. No state information, but TCL properties are known	yes, through analytical derivation	yes

Table 2: Options for control and information transfer.

From the structure of \mathbf{B} it can be seen that no probability mass (or TCLs) can vanish from the state space through control actions. Based on our sign convention, positive elements of $\mathbf{u}_{\text{bin},i}$ correspond to turning TCLs ON and negative elements correspond to turning TCLs OFF. This property can be exploited by constraining \mathbf{u}_{bin} to only non-negative or non-positive values, if the system is designed such that the control can only work in one direction.

We target TCLs according to B3 as follows: If the controller chooses to switch a certain number of loads in a bin, it converts this number to a switch probability using available state information. Then, loads (using a local random number generator) switch with this probability, therefore the same control signal causes differing responses within the TCL population. This avoids the need to directly address individual TCLs, thus reducing the communication effort between the central controller and the TCL population.

3.2.2 TCL Population \rightarrow Controller

Knowing the measured temperatures of the entire population is the most convenient condition for control (Option C1 in Table 2) because this means that the \mathbf{A} -matrix can be identified using Markov Chain techniques, as in Section 2.2.2. However, this option is the most expensive to implement, so we would like to reduce the need for measured state information. Unfortunately, it is not beneficial to measure a subset of the states. Individual TCLs traverse all states, so measuring a subset of states would imply that all TCLs would have to be instrumented, which is equivalent to having full state information¹. We propose instead to instrument a subset of TCLs (Option C2). At intervals, each TCL will transmit its internal temperature, with respect to its dead-band, to the central controller. Aggregating these results, we have a noisy measurement of the state bin vector, which can be used directly or together with an observer (e.g., Kalman filter). The \mathbf{A} -matrix can be identified using Markov Chain techniques on the noisy data.

Integrating measured information into the control system boils down to designing the output equation of the state bin transition model. The measured output $\mathbf{y} = \mathbf{C}\mathbf{x} + \mathbf{v}$, where \mathbf{v} is measurement noise, does not include a direct feed-through term $\mathbf{D}\mathbf{u}$. This is because the con-

¹A possible exception to this is TCLs reporting their switching actions to the controller, which yields information about the bins at the boundary of the dead-band.

troller acts on the TCLs, which translates to a difference in output one time step later. The matrix \mathbf{C} reflects whether only aggregate power is measured, or a measurement of the full state vector is included. In the former case:

$$\mathbf{C} = \underbrace{P_{\text{rate}} N_{\text{TCL}}}_{:=C_p} \left[\underbrace{0, \dots, 0}_{\frac{N_{\text{bin}}}{2}}, \underbrace{1, \dots, 1}_{\frac{N_{\text{bin}}}{2}} \right]. \quad (16)$$

In the latter case, the matrix \mathbf{C} becomes:

$$\mathbf{C} = \begin{bmatrix} \mathbf{I}(N_{\text{bin}}) \\ 0, \dots, 0, C_p, \dots, C_p \end{bmatrix}, \quad (17)$$

where $\mathbf{I}(N_{\text{bin}})$ is the N_{bin} -dimensional identity matrix and the vector $\mathbf{v} = [v_1, \dots, v_{N_{\text{bin}}}, v_p]^T$ is measurement noise resulting from partial instrumentation of the population and other sources.

3.3 Controllability and Observability Properties

The pair $[\mathbf{A}, \mathbf{B}]$ is not controllable: the controllability matrix is of rank $n - 1$. One degree of freedom is lost because the controller cannot drive all states to zero since the states, which represent the fraction of TCLs in each bin, must sum to 1. However, if \mathbf{C} is defined as in (16), the system output is output controllable and aggregate power can be tracked. If \mathbf{C} is defined as in (17), the system is not output controllable: the output controllability matrix is of rank $n - 2$. One degree of freedom is lost because the controller cannot drive all states to zero and another is lost because aggregate power is dependent on the states. However, aggregate power can still be tracked. The pair $[\mathbf{A}, \mathbf{C}]$ is observable for both options of \mathbf{C} .

4 Control Approach

4.1 Initial Considerations

The state bin transition model is linear because the control input influences an *absolute* (albeit approximate) number of TCLs in a certain bin in each time step. This is a significant difference to the formulation in [7], in which the controller influences a *percentage* of TCLs in a bin, resulting in a nonlinear model. In our formulation, the system can theoretically be controlled to have less than zero or more than all of the TCLs in a bin, ($x_i < 0$ or $x_i > 1$ for any $i = 1, \dots, N_{\text{bin}}$), which is physically impossible. Consequently, we choose to use MPC because it can incorporate inequality constraints on states, not just inputs.

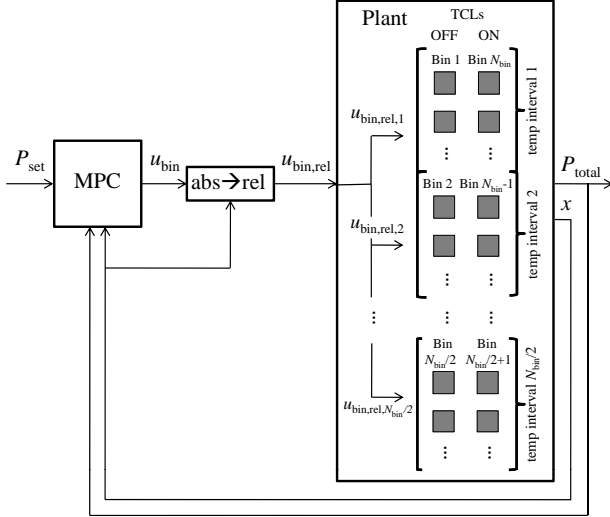


Figure 3: Control structure for TCL population.

MPC schemes are usually applied to control problems where certain variables must be controlled to (possibly varying) setpoints taking into account the (usually slow) system dynamics, such as in the process industry. In this case, special attention has to be given to the sampling time and prediction horizon in order to represent the dynamics properly within the controller. In our case, switching actions imposed on TCLs occur (in principle) instantly when the control is applied. Due to the state bin transition model formulation, this happens one step after the control is applied, but the whole range of possible outputs can be traversed in this one step. This means that a very short prediction horizon (minimum 2 steps for implementation reasons) is sufficient to capture the relevant input-output dynamics. In cases where the target trajectory is known several steps in advance, a longer prediction horizon can provide optimization potential for minimizing the control impact on the TCL population.

4.2 Control Problem Formulation

Figure 3 illustrates the control structure consistent with the control and information transfer choices described in the previous section. The MPC outputs the control signal u_{bin} (in terms of probability mass), which is converted to $u_{\text{bin,rel}}$ (in terms of switch probabilities) by using state information. Note that this preserves linearity in the plant / control input model. To formulate the control problem, a set of constraints, defined by the properties of probability mass evolving in the state space, is introduced:

$$0 \leq x_i \leq 1 \quad \forall \quad i = 1, \dots, N_{\text{bin}} \quad , \quad (18)$$

$$-1 \leq u_{\text{bin},i} \leq 1 \quad \forall \quad i = 1, \dots, N_{\text{bin}}/2 \quad . \quad (19)$$

Beside these principal physical constraints, other constraints may be imposed based on the desired behavior of the controller.

A number of suitable cost function designs achieving certain design goals are discussed below. Here, we restrict ourselves to quadratic cost functions, as these usually yield smooth control behavior and are easy to solve

by standard quadratic programming techniques. In the choice of penalty factors, there are several options, which include:

- penalization of the tracking error, $P_{\text{total}} - P_{\text{set}}$,
- penalization of the control input vector, u_{bin} , and
- penalization of the state vector deviation from a desired value, $x - x_{\text{set}}$.

In its most general form, the cost function is:

$$J_k = \sum_{l=k}^{k+N_{\text{pred}}-1} (Q_{\text{track}} (P_{\text{total},l} - P_{\text{set},l})^2 + u_{\text{bin},l}^T R_{\text{bin}} u_{\text{bin},l} + (x_l - x_{\text{set},l})^T Q (x_l - x_{\text{set},l})) \quad (20)$$

The following options are considered for influencing the control behavior in the desired way:

- $R_{\text{bin}} = \mathbf{0}$, $Q = \mathbf{0}$ achieves the minimization of setpoint tracking error without any attention to the control action.
- To balance control actions against the tracking error, R_{bin} should be non-zero. In the most simple form, all main diagonal elements are chosen positive and equal, resulting in a uniform evolution of all elements of u_{bin} . A reduction of switching actions can be achieved by concentrating the control actions on bins with a higher index number, which causes TCLs that are closer to switching to be switched preferably. This can be achieved by making R_{bin} a diagonal matrix with a monotonically decreasing series of elements on the main diagonal. Depending on the desired penalization, the series can decrease linearly, quadratically, or exponentially.
- Another alternative cost function design takes advantage of the flexibility provided by the bin transition model and results in a more complex controller behavior, although it is easy to formulate. The state vector x can be penalized by the diagonal matrix Q with respect to its deviation from a certain desired probability mass profile (e.g., the steady state profile) among the bins.

5 Case Study

5.1 System setup and scenarios

To evaluate the performance of the control system, a simulated population of individual TCLs with process noise set to zero ($w_{i,t} = 0$), is controlled to track a highly variable setpoint. The test trajectory is composed of four different sections: (1) a series of steps resembling energy market or tertiary frequency control dispatch, (2) a sinusoid representing a smooth power adjustment trajectory,

Options	Parameters			
	N_{pred}	Q_{track}	Q	R_{bin}
cost. u_{bin}	2	10^3	$\mathbf{0}$	$\text{diag}([1, \dots, 1])$
dec. u_{bin}	2	10^3	$\mathbf{0}$	$*R_{\text{bin,dec}}$
state	2	10^3	$\text{diag}([1, \dots, 1])$	$\mathbf{0}$

$$*R_{\text{bin,dec}} = \underbrace{[10^3, \dots, 10^3]}_{N_{\text{bin}}/2}, \underbrace{[N_{\text{bin}}/2, \dots, 1]}_{N_{\text{bin}}/2}$$

Table 3: Options for controller parameterization.

	Case 1	Case 2	Case 3	Case 4	Case 5	Case 6	Case 7	Case 8	Case 9	Case 10	Case 11	Case 12
Parameter set												
N_{bin}	40	40	40	40	40	60	60	60	80	80	80	80
Control penalty (see Table 3)	cst. u_{bin}	cst. u_{bin}	cst. u_{bin}	dec. u_{bin}	state	cst. u_{bin}	dec. u_{bin}	state	cst. u_{bin}	dec. u_{bin}	state	state
Percentage measured [%]	100	30	10	100	100	100	100	100	100	100	100	10
Performance indicators												
Ctrl. error (rel. RMSE) [%]	1.87	1.79	2.27	2.25	1.89	1.03	1.68	1.03	0.80	1.39	0.80	1.32
Mean switching increase [%]	168	201	294	258	246	193	301	280	185	317	292	343
Max switching increase [%]	364	375	525	725	475	400	767	645	367	845	509	633
Min switching increase [%]	44	93	138	38	112	82	46	119	80	38	129	150

Table 4: Performance comparison.

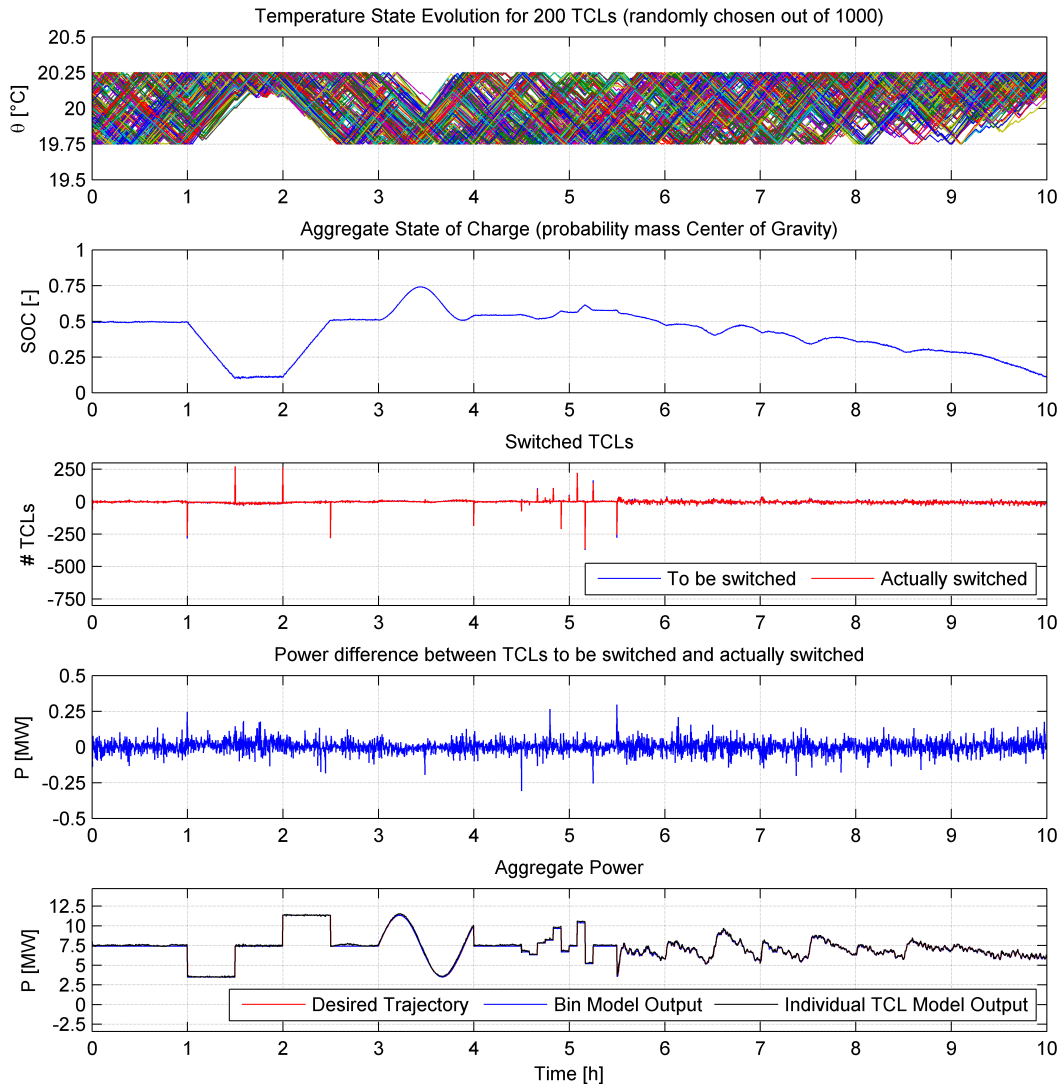


Figure 4: Simulation result for full state information.

(3) some 5-minute power changes which is relevant for short-term ancillary service markets [8], and (4) a real Load Frequency Control signal obtained from a European Transmission System Operator. The choice of these diverse setpoints demonstrates the versatility of the control approach.

The simulated model is parameterized according to the discussion in Section 2. Specifically, devices are heterogeneous in thermal capacitance (C) but homogeneous with respect to rated power (P_{rate}) and thermal resistance (R). Therefore, the devices are heterogeneous in the TCL parameter a and homogeneous in the temperature gain θ_g (2). A real population of TCLs would be heterogeneous in θ_g ,

which would lead to additional diversity in dynamics during the ON phase. We will study the impact of rated power and thermal resistance heterogeneity on model outcomes in future work.

Numerical parameters are shown in Table 1. Table 3 presents options for parameterizing the controller. Specifically, the columns represent three different choices for penalizing control inputs: “cst. u_{bin} ” denotes a constant (equal) penalization of the elements in u_{bin} , “dec. u_{bin} ” means that the second half of u_{bin} is penalized in a (linearly) decreasing way, and “state” means that the state is penalized with respect to the steady state of the bin transition model.

We define a set of simulation scenarios that provide insight into the control behavior for different parameterizations of the controller. One scenario will be shown graphically, while the others will only be considered in numerical performance comparisons. Twenty-seven scenarios can be explored by permutating the following:

- Variation of N_{bin} : 40, 60, 80 state bins
- Variation of cost function design:
 - equal penalty on u_{bin} for all elements
 - high penalty for first half of u_{bin} , lower and linearly decreasing penalty for the second half
 - state penalty
- Variation of state information available to the controller: 100%, 30%, 10%

5.2 Numerical Results

Numerical simulations were carried out using the scenario definitions above. For the Model Predictive Control setup, the MATLAB toolbox YALMIP [9] was used. Figure 4 shows a simulation of the controlled system with full state vector information available to the controller. The parameters correspond to Case 8 in Table 4. The desired trajectory is tracked with good performance.

The state of charge (SOC) of the aggregate system is defined as the center of gravity of device states with respect to the temperature dead-band. For cooling TCLs, all TCLs concentrated at the upper dead-band corresponds to $SOC = 0$, at the lower dead-band to $SOC = 1$. The opposite relation holds for heating TCLs. Information on the aggregate SOC is relevant to power system operators, allowing them to dispatch the TCL population like an energy storage device.

Table 4 shows a numerical performance comparison of 12 of the 27 simulated scenarios. The parameter sets are presented in the upper section of the table, whereas the performance indicators are shown below. It can be seen that the relative Root Mean Square Error (RMSE, normalized by average aggregate power) of the setpoint tracking lies in the range of 0.8 – 2.27% depending on the simulated case. These results suggest that tracking can be improved with a larger number of bins and / or a larger set of directly measured TCLs, but the effect is relatively small.

Furthermore, we investigate how the control algorithm influences the ON/OFF switching frequency of the TCL population. The number of ON/OFF alterations (switchings) in each simulated case is compared with the number of switchings in the case where all TCLs operate autonomously. The average increase in switching frequency is roughly 170 to 300%. This significant increase is due to the shape of the desired trajectory which imposes large deviations of active power consumption with respect to the absence of control, forcing the population to sometimes operate far away from the mean SOC of 0.5.

6 Conclusion and Outlook

This paper demonstrates the ability of a large number of thermostatically controlled loads to track a variable power

signal. The bin transition modeling technique provides a model of the aggregate dynamic behavior of the TCL population which is accurate enough for use with model-based control techniques. Heterogeneity of the population was incorporated into the approach and tracking performance was demonstrated with a set of 1,000 TCLs.

Further research will address the reduction of communication requirements so as to increase the chance for cost-effective practical application of the technique. An issue is the quality of available state information to the controller, which makes the application of state estimation and filtering techniques promising for increasing tracking performance in communication-constrained environments.

Acknowledgements

We thank Taylor Keep, Froy Sifuentes, Dave Auslander and especially Hosam Fathy for stimulating conversations and feedback. The first author would like to thank Duncan Callaway for the excellent opportunity to conduct research within his group at UC Berkeley, and *swisselectric research* for financial support of the project *Local Load Management*. The second author is funded by a UC Berkeley Chancellor’s Fellowship.

REFERENCES

- [1] Y.V. Makarov, C. Loutan, Jian Ma, and P. de Mello. Operational impacts of wind generation on California power systems. *IEEE Transactions on Power Systems*, 24(2):1039-1050, 2009.
- [2] J. Taneja, D. Culler, and P. Dutta. Towards cooperative grids: Sensor/actuator networks for renewables integration. In *Proceedings of IEEE SmartGridComm*, Gaithersburg, MD, USA, 2010.
- [3] N. Lu, D.P. Chassin, and S.E. Widergren. Modeling uncertainties in aggregated thermostatically controlled loads using a state queueing model. *IEEE Transactions on Power Systems*, 20(2):725-733, 2005.
- [4] S. Koch, M. Zima, and G. Andersson. Active coordination of thermal household appliances for load management purposes. In *Proceedings of IFAC Symposium on Power Plants and Power Systems Control*, Tampere, Finland, 2009.
- [5] S. Koch, F. Soto Barcenás, and G. Andersson. Using controllable thermal household appliances for wind forecast error reduction. In *Proceedings of IFAC Conference on Control Methodologies and Technology for Energy Efficiency*, Vilamoura, Portugal, 2010.
- [6] D.S. Callaway. Tapping the energy storage potential in electric loads to deliver load following and regulation, with application to wind energy. *Energy Conversion and Management*, 50(9):1389-1400, 2009.
- [7] S. Bashash and H.K. Fathy. Modeling and Control Insights into Demand-side Energy Management through Setpoint Control of Thermostatic Loads. In *Proceedings of American Control Conference*, San Francisco, CA, USA, 2011.
- [8] J. Huang, P. Yalla, and T. Yong. New real time market applications at the California Independent System Operator (CAISO). In *Proceedings of IEEE PES Power Systems Conference and Exposition*, New York, NY, USA, 2004.
- [9] J. Löfberg, Yalmip: A toolbox for modeling and optimization in MATLAB, In *Proceedings of the CACSD Conference*, Taipei, Taiwan, 2004. [Online]. Available: <http://users.isy.liu.se/johanl/yalmip>